arXiv:1409.7419v1 [stat.ME] 25 Sep 2014

Identifying the number of clusters in discrete mixture models

Cláudia Silvestre,[a] * Margarida G. M. S. Cardoso,[b] Mário A. T. Figueiredo[c]

[a] *Escola Superior de Comunicação Social, Instituto Politécnico de Lisboa, Portugal*;

[b] *Department of Quantitative Methods, ISCTE - Lisbon University Institute, Portugal*;

[c] *Instituto de Telecomunicações, Instituto Superior Técnico, Universidade de Lisboa, Portugal,*

September 29, 2014

Abstract

Research on cluster analysis for categorical data continues to develop, with new clustering algorithms being proposed. However, in this context, the determination of the number of clusters is rarely addressed. In this paper, we propose a new approach in which clustering of categorical data and the estimation of the number of clusters is carried out simultaneously. Assuming that the data originate from a finite mixture of multinomial distributions, we develop a method to select the number of mixture components based on a *minimum message length* (MML) criterion and implement a new *expectation-maximization* (EM) algorithm to estimate all the model parameters. The proposed EM-MML approach, rather than selecting one among a set of pre-estimated candidate models (which requires running EM several times), seamlessly integrates estimation and model selection in a single algorithm. The performance of the proposed approach is compared with other well-known criteria (such as the *Bayesian information criterion*–BIC), resorting to synthetic data and to two real applications from the European Social Survey. The EM-MML computation time is a clear advantage of the proposed method. Also, the real data solutions are much more parsimonious than the solutions provided by competing methods, which reduces the risk of model order overestimation and increases interpretability.

keywords: finite mixture model; EM algorithm; model selection, minimum message length; categorical data

1 Introduction

Clustering is a technique commonly used in several research and application areas, such as social sciences, medicine, biology, engineering, computer science, image analysis, bioinformatics, and marketing. The goal of clustering

is to discover or uncover groups in data. To this end, there are essentially two different approaches: distance-based, where a distance or a similarity measure between objects is defined and similar objects are assigned to the same group; model-based, where the data are assumed to be generated by a finite mixture model, and objects are assigned to groups based on the corresponding estimates of the posterior probabilities [1].

Most of the clustering techniques are focused on numerical data and can not be applied directly to categorical data. In fact, clustering techniques for categorical data are more challenging [2], and there are fewer techniques available [3]. When using distance-based clustering approaches, one needs to resort to specific similarity measures to deal with categorical features – e.g [4]. In this context, the determination of the number of clusters is commonly based on clustering quality measures and the corresponding graphical displays, such as dendrograms (when using hierarchical clustering techniques) or entropy related graphics [5]. In model-based approaches for numerical data, the usual choice is a mixture of Gaussians (e.g. [6], [7], [8]); when referring to categorical data, a mixture of multinomials – discrete mixture model – is usually considered (e.g. [3], [9], [10], [11]). In order to determine the number of groups in discrete mixture models, information criteria are commonly used: e.g., the *Bayesian information criterion* (BIC) [12] or the *Akaike information criterion* (AIC) [13]. These criteria look for a balance between the model's fit to the data (which ocrresponds to maximizing the likelihood function) and parsimony (using penalties associated with measures of model complexity), thus trying to avoid over-fitting. The use of information criteria follows the estimation of candidate finite mixture models for which a predetermined number of groups is indicated, generally resorting to an EM (*expectation-maximization*) algorithm, [14].

In this work, we focus on determining the number of groups while clustering categorical data, using an EM embedded approach to estimate the number of groups. The novelty of this approach is that it does not rely on selecting among a set of pre-estimated candidate models, but rather integrates estimation and model selection in a single algorithm. We capitalize on the approach developed by Figueiredo and Jain [15] for clustering continuous data and extend it for dealing with categorical data. The proposed method is based on a *minimum message length* (MML) criterion to select the number of clusters and on an EM algorithm to estimate the model parameters; our implementation follows a previous version described in [16].

The paper is organized as follows: Section 2 reviews the finite mixture model-based approach for clustering and addresses the case of categorical data; Section 3 provides an introduction to the topic of model selection for discrete finite mixtures; in Section 4 we describe the proposed EM-MML based algorithm; in Section 5, the experimental results, based on synthetic and real data, illustrate the performance of the EM-MML approach. Concluding remarks are summarized in Section 6.

2 Clustering with finite mixture models

Finite mixture models offer a model-based approach to clustering, exhibiting some competitive advantages when compared to alternative methods: besides producing the allocation of observations to clusters, they yield estimates of within-clusters joint probability functions for the base variables; moreover, they provide means to allocate new observations to groups; finally, when used with information criteria, mixture models provide a statistical framework to determine the number of clusters.

The literature on finite mixture models and their application is vast, including some books covering theory, geometry, and applications [17], [18], [19], [20], [21], [22]. For example, in market segmentation, cluster analysis via finite mixture models has replaced more traditional cluster analysis, such as K-Means algorithm, as the state of the art [23].

Finite mixture models have played an important role in the history of modern statistics. One of the first applications of mixture models is due to Newcomb [24], who used a mixture of Gaussians to analyse a collection of observations of transits of Mercury. Pearson [25] fitted a mixture of two Gaussians with unequal variances in an analysis of different species of crabs. Among many other examples of applying mixture models, MacDonald and Pitcher [26] analysed single species of fish in a lake, using a mixture of Gaussians where each component consists of the fish of a single yearly spawning of that species. Another example is given by Do and McLachlan [27], where mixture of Gaussians was used to study the populations of rats that were being eaten by a particular species of owl, with the distinct rat species corresponding to the components of mixtures. Al-Hussaini and Abdel-Hamid [28] studied the behavior of failure time of a device; they fitted a mixture of components, each of which represents a different cause of failure. In their research, a special attention was paid to mixtures of two Weibull components, but two exponential components, two Rayleigh components, and mixture of Rayleigh and Weibull components, were also analysed. In these examples, the segmentation process uncovers physically meaningful components.

When applying finite mixture models to social sciences, the analyst may be confronted with the need to uncover sub-populations based on qualitative indicators. In this context, the use of mixture models of categorical data is particularly pertinent. For example, for clustering categorical data from multiple choice questions, a mixture of multinomial distributions is used in order to market products [29].

2.1 Definitions and concepts

Let $\mathbf{Y} = \{\underline{y}_i,\ i = 1, \ldots, n\}$ be a set of n independent and identically distributed (i.i.d.) sample observations of a random vector, $\underline{Y} = [Y_1, \ldots, Y_L]'$. If $\underline{Y}$ follows a mixture of K components densities, $f(\underline{y}|\underline{\theta}_k)$ $(k = 1, \ldots, K)$,

with probabilities $\{\alpha_1, \ldots, \alpha_K\}$, the probability (density) function of $\underline{Y}$ is

$$f(\underline{y}|\Theta) = \sum_{k=1}^{K} \alpha_k f(\underline{y}|\underline{\theta}_k),$$

where $\Theta = \{\underline{\theta}_1, \ldots, \underline{\theta}_K, \alpha_1, \ldots, \alpha_K\}$ is the set of all the parameters of the model and $\underline{\theta}_k$ are the distributional parameters defining the k-th component. The proportions, also called *mixing probabilities*, are subject to the usual constraints: $\sum_{k=1}^{K} \alpha_k = 1$ and $\alpha_k \geq 0$, $k = 1, \ldots, K$.

The log-likelihood of the observed set of sample observations is

$$\log f(\mathbf{Y}|\Theta) = \log \prod_{i=1}^{n} f(\underline{y}_i|\Theta) = \sum_{i=1}^{n} \log \sum_{k=1}^{K} \alpha_k f(\underline{y}_i|\underline{\theta}_k).$$

In clustering, the identity of the component that generated each sample observation is unknown. The observed data $\mathbf{Y}$ is therefore regarded as incomplete, where the missing data is a set of indicator variables $\mathbf{Z} = \{\underline{z}_1, ..., \underline{z}_n\}$, each taking the form $\underline{z}_i = [z_{i1}, ..., z_{iK}]'$, where z_{ik} is a binary indicator: z_{ik} takes the value 1 if the observation $\underline{y}_i$ was generated by the k-th component, and 0 otherwise. It is usually assumed that the $\{\underline{z}_i,\ i = 1, \ldots, n\}$ are i.i.d., following a multinomial distribution of K categories, with probabilities $\{\alpha_1, \ldots, \alpha_K\}$. The log-likelihood of complete data $\{\mathbf{Y}, \mathbf{Z}\}$ is given by

$$\log f(\mathbf{Y}, \mathbf{Z}|\Theta) = \sum_{i=1}^{n} \sum_{k=1}^{K} z_{ik} \log \left[\alpha_k f(\underline{y}_i|\underline{\theta}_k)\right].$$

2.2 Discrete Finite Mixture Models

Consider that each variable in $\underline{Y}$, Y_l $(l = 1, \ldots, L)$ can take one of C_l categories. Conditionally on having been generated by the k-th component of the mixture, each Y_l is thus modeled by a multinomial distribution with n_l trials, C_l categories, and non-negative parameters $\underline{\theta}_{kl} = \{\theta_{klc},\ c = 1, \ldots, C_l\}$, with $\sum_{c=1}^{C_l} \theta_{klc} = 1$. For a sample $y_{il} (i = 1, \ldots, n)$ of Y_l, we denote as y_{ilc} the number of outcomes in category c, which is a sufficient statistic; naturally, $\sum_{c=1}^{C_l} y_{ilc} = n_l$. Thus, with $\underline{\theta}_k = \{\underline{\theta}_{k1}, \ldots, \underline{\theta}_{kL}\}$ and $\Theta = \{\underline{\theta}_1, \ldots, \underline{\theta}_K, \alpha_1, \ldots, \alpha_k\}$, the log-likelihood function, for a set of observations corresponding to a discrete finite mixture model (mixture of multinomials) is

$$\log p(\mathbf{Y}|\Theta) = \sum_{i=1}^{n} \log \sum_{k=1}^{K} \alpha_k \prod_{l=1}^{L} \left[n_l! \prod_{c=1}^{C_l} \frac{(\theta_{klc})^{y_{ilc}}}{y_{ilc}!} \right]. \tag{1}$$

In order to estimate the parameters of this mixture, it is necessary to ensure that they are identifiable. In particular, in the case of a mixture of

multinomials, specific identifiability condition must be fulfilled: a mixture of multinomial distributions is identifiable if $T \geq 2K - 1$, where T is the number of trials of each multinomial distribution (see details in [30] and [31]).

As in the Gaussian case, the log-likelihood in (1) can be seen as corresponding to a missing-data problem, where the missing data has exactly the same meaning and structure as above. The log-likelihood of the complete data $\{\mathbf{Y}, \mathbf{Z}\}$ is thus given by

$$\log p(\mathbf{Y}, \mathbf{Z}|\Theta) = \sum_{i=1}^{n} \sum_{k=1}^{K} z_{ik} \log \left(\alpha_k \prod_{l=1}^{L} \left[n_l! \prod_{c=1}^{C_l} \frac{(\theta_{klc})^{y_{ilc}}}{y_{ilc}!} \right] \right). \tag{2}$$

2.3 Estimation of finite mixture models via the EM algorithm

To obtain a *maximum-likelihood* (ML) or *maximum a posteriori* (MAP) estimate of the parameters of a multinomial mixture, the well-known EM algorithm is usually the tool of choice ([14] and [32]). EM is an iterative algorithm, which alternates between two steps, the *expectation step* (E-step) and the *maximization step* (M-step), described next.

E-step: Compute the expectation of the complete log-likelihood (2), with respect to the missing variables $\mathbf{Z}$, given the observed data $\mathbf{Y}$, and the current parameter estimate $\widehat{\Theta}^{(t)}$ (where t is the iteration counter). Because $\log p(\mathbf{Y}, \mathbf{Z}|\Theta)$ is linear with respect to $\mathbf{Z}$ (as is clear in (2)),

$$\mathbb{E}\left[\log p(\mathbf{Y}, \mathbf{Z}|\Theta) \middle| \mathbf{Y}, \widehat{\Theta}^{(t)} \right] = \log p(\mathbf{Y}, \bar{\mathbf{Z}}^{(t)} \middle| \Theta)$$

where each element $\bar{z}_{ik}^{(t)}$ of $\bar{\mathbf{Z}}^{(t)}$ is given by

$$\bar{z}_{ik}^{(t)} = \mathbb{E}\left[Z_{ik} \middle| \mathbf{Y}, \widehat{\Theta}^{(t)} \right] = \mathbb{P}\left[Z_{ik} = 1 \middle| \underline{y}_i, \widehat{\Theta}^{(t)} \right] = \frac{\alpha_k \; f(\underline{y}_i|\underline{\theta}_k)}{\sum_{j=1}^{K} \alpha_j \; f(\underline{y}_i|\underline{\theta}_j)}, \tag{3}$$

since Z_{ik} is binary and conditionally independent from all $\underline{y}_j$, for $j \neq i$; the third equality results simply from Bayes' law.

M-step: Update the parameter estimates by maximizing the current estimate of the expected complete log-likelihood

$$\widehat{\Theta}^{(t+1)} = \arg\max_{\Theta} \; \log p(\mathbf{Y}, \bar{\mathbf{Z}}^{(t)} \middle| \Theta) + \log p(\Theta), \tag{4}$$

where $p(\Theta)$ is a prior, in the case of MAP estimation; in the case of ML estimation, the term $\log p(\Theta)$ is absent.

A more detailed derivation of the EM algorithm, its convergence properties, extensions, and instances for different type of missing-data models, can be found, e.g., in [33], [34].

3 Model selection for categorical data

Model selection is an important problem in statistical analysis [35]. In model-based clustering, the term *model selection* usually refers to the problem of determining the number of clusters, although it may also refer to the problem of selecting the structure of the clusters. Model-based clustering provides a statistical framework to solve this problem [36], usually resorting to *information criteria*. The rationale of such criteria is that fitting a model with a large number of clusters requires estimation of a very large number of parameters and a consequent loss of precision in these estimates. Therefore, one should penalize excessive model complexity and simultaneously try to increase the model's fit to the data, based on the likelihood function. The best-known information criteria are BIC, AIC, and their modifications, namely the *consistent AIC* (CAIC) [37], and the *modified AIC* (MAIC) [38]. Other criteria that have been proposed include the *integrated completed likelihood* (ICL) [39], the *minimum description length* (MDL) [40], and the *minimum message length* (MML) criterion [41]. Information criteria are well-known and easily implemented, the final model being selected according to a compromise between its fit to data and its complexity.

In this work, we use the MML criterion to choose the number of components of a mixture of multinomials. MML is based on the information-theoretic view of estimation and model selection, according to which an adequate model is one that allows a short description of the observations [41]. MML-type criteria evaluate statistical models according to their ability to compress a message containing the data, looking for a balance between choosing a simple model and one that describes the data well.

According to Shannon's information theory, if Y is some random variable with probability distribution $p(y|\Theta)$, the optimal code-length (in an expected value sense) for an outcome y is $l(y|\Theta) = -\log_2 p(y|\Theta)$, measured in bits (from the base-2 logarithm) [42]. If Θ is unknown, the total code-length function has two parts: $l(y, \Theta) = l(y|\Theta) + l(\Theta)$; the first part encodes the outcome y, while the second part encodes the parameters of the model. The first part corresponds the fit of the model to the data (better fit corresponds to higher compression), while the second part represents the complexity of the model. Different ways to compute $l(\Theta)$ are derived from different statistical frameworks and yield different flavors of information criteria, namely MDL and MML, where MML admits the existence of a prior $p(\Theta)$, while MDL does not.

The message length function for a mixture of distributions (as developed in [43]) is:

$$l(y,\Theta) = -\log p(\Theta) - \log p(y|\Theta) + \frac{1}{2}\log|I(\Theta)| + \frac{C}{2}\left(1 - \log(12)\right), \quad (5)$$

where $p(y|\Theta)$ is the likelihood function, $I(\Theta) \equiv -E\left[\frac{\partial^2}{\partial\Theta^2}\log p(Y|\Theta)\right]$ is the expected Fisher information matrix, $|I(\Theta)|$ its determinant, and C is the the number of parameters of the model that need to be estimated. For example, for the K mixture multinomial distributions presented in (1),

$$C = (K-1) + K\left(\sum_{l=1}^{L}(C_l - 1)\right).$$

The expected Fisher information matrix of a mixture leads to a complex analytical form of MML which cannot be easily computed. To overcome this difficulty, Figueiredo and Jain [15] replace the expected Fisher information matrix by its complete-data counterpart $I_c(\Theta) \equiv -E\left[\frac{\partial^2}{\partial\theta^2}\log p(Y,Z|\Theta)\right]$. Also, they adopt independent Jeffreys' *priors* for the mixture parameters. The resulting message length function is

$$l(y,\Theta) = \frac{M}{2}\sum_{k:\,\alpha_k>0}\log\left(\frac{n\,\alpha_k}{12}\right) + \frac{k_{nz}}{2}\log\frac{n}{12} + \frac{k_{nz}(M+1)}{2} - \log p(y,\Theta) \quad (6)$$

where M is the number of parameters specifying each component (the dimension of each $\underline{\theta}_k$) and k_{nz} the number of components with non zero probability (for more details on the derivation of (6), see [15], [43]).

4 The proposed MML based EM algorithm

In order to estimate a mixture of multinomials, we propose to use a variant of the EM algorithm (herein termed EM-MML), which integrates both estimation and model selection, by directly minimizing (6). This algorithm is an extension, to the multinomial case, of the approach developed by in [15] for clustering continuous data, based on a Gaussian mixture model.

The algorithm results from observing that (6) contains, in addition to the log-likelihood term, an explicit penalty on the number of components (the two terms proportional to k_{nz}), and a term (the first one) that can be seen as a log-prior on the α_k parameters of Θ, that will directly affect the M-step. Finally, notice that in the presence of a set of multinomial observations $\mathbf{Y} = \{\underline{y}_i,\ i = 1,\ldots,n\}$, the log-likelihood $\log p(\mathbf{Y}\big|\Theta)$ is as given in (1).

E-step: The E-step of the EM-MML is precisely the same as in the case of ML or MAP estimation, since the generative model for the data is

the same. Since we are dealing with a multinomial mixture, we simply have to plug the corresponding multinomial probability function in (3), yielding

$$\bar{z}_{ik}^{(t)} = \frac{\alpha_k \prod_{l=1}^{L} \left[n_l! \prod_{c=1}^{C_l} \frac{(\widehat{\theta}_{klc}^{(t)})^{y_{ilc}}}{y_{ilc}!} \right]}{\sum_{j=1}^{K} \alpha_j \prod_{l=1}^{L} \left[n_l! \prod_{c=1}^{C_l} \frac{(\widehat{\theta}_{jlc}^{(t)})^{y_{ilc}}}{y_{ilc}!} \right]}, \tag{7}$$

for $i = 1, \ldots, n$ and $k = 1, \ldots, K$.

M-step: For the M-step, noticing that the first term in (6) can be seen as the negative log-prior $-\log p(\alpha_k) = \frac{C-K+1}{2K} \log \alpha_k$ (plus a constant), and enforcing the conditions that $\alpha_k \geq 0$, for $k = 1, ..., K$ and that $\sum_{k=1}^{K} \alpha_k = 1$, yields the following updates for the estimates of the α_k parameters:

$$\widehat{\alpha}_k^{(t+1)} = \frac{\max\left\{0, \sum_{i=1}^{n} \bar{z}_{ik}^{(t)} - \frac{C-K+1}{2K}\right\}}{\sum_{j=1}^{K} \max\left\{0, \sum_{i=1}^{n} \bar{z}_{ij}^{(t)} - \frac{C-K+1}{2K}\right\}}, \tag{8}$$

for $k = 1, ..., K$. Notice that, some $\widehat{\alpha}_k^{(t+1)}$ may be zero; in that case, the k-th component is excluded from the mixture model. The multinomial parameters corresponding to components with $\widehat{\alpha}_k^{(t+1)} = 0$ need not be further calculated, since these components do not contribute to the likelihood. For the components with non-zero probability, $\widehat{\alpha}_k^{(t+1)} > 0$, the estimates of multinomial parameters are updated to their standard weighted ML estimates:

$$\widehat{\theta}_{klc}^{(t+1)} = \frac{\sum_{i=1}^{n} \bar{z}_{ik}^{(t)} y_{ilc}}{n_l \sum_{i=1}^{n} \bar{z}_{ik}^{(t)}}, \tag{9}$$

for $k = 1, \ldots, K$, $l = 1, \ldots, L$, and $c = 1, \ldots, C_l$. Notice that, in accordance with the meaning of the θ_{klc} parameters, $\sum_{c=1}^{C_l} \widehat{\theta}_{klc}^{(t+1)} = 1$.

The EM-MML algorithm for clustering categorical data and selecting the number of clusters simultaneously is summarized in Figure 1.

Figure 1: The EM-MML algorithm

Input: data: $\mathbf{Y} = \{\underline{y}_i,\ i = 1, \ldots, n\}$
the minimum number of segments: K_{min}
the maximum number of segments: K_{max}
minimum increasing threshold for the log-likelihood function: δ
Ouput: number of segments: K
segments probabilities: $\{\alpha_1, \ldots, \alpha_K\}$
multinomial parameters:$\{\underline{\theta}_1, \ldots, \underline{\theta}_K\}$

Initialization: initialization of the parameters resorts to the empirical distribution:
$p(\underline{y}_l|\underline{\theta}_{lk})$, $(l = 1, \ldots, L\ ;\ k = 1, \ldots, K_{max})$
set the segment's probability: $\alpha_k = 1/K_{max}$ $(k = 1, \ldots, K_{max})$
store the initial log-likelihood
store the initial message length (iml)
$minml \leftarrow iml$

continue $\leftarrow 1$
while continue $= 1$ **do**
  while increases on log-likelihood are above δ **do**
    $k = K_{max}$
    while $k \geqslant K_{min}$ **do**
      compute $\widehat{\alpha}_k$ according to (8)
      if $\widehat{\alpha}_k = 0$
        remove the component k of the mixture
        $K_{max} \leftarrow K_{max} - 1$
        $\{\widehat{\alpha}_1, \ldots, \widehat{\alpha}_{K_{max}}\} \leftarrow \left\{ \frac{\widehat{\alpha}_1}{\sum_{k=1}^{k_{max}} \widehat{\alpha}_k}, \ldots, \frac{\widehat{\alpha}_{K_{max}}}{\sum_{k=1}^{k_{max}} \widehat{\alpha}_k} \right\}$
      else
        compute $\widehat{\theta}$ according to (9)
        E-step according to (7)
        $k \leftarrow k - 1$
      end if
    end while
    compute log-likelihood
    compute message length (ml)
  end while
  if $ml < minml$
    $minml \leftarrow ml$
    segment's probability $\leftarrow$ current parameters
    multinomial probabilities $\leftarrow$ current parameters
  end if
  if $K_{max} > K_{min}$
    $k_{remove} = argmin\,\{\alpha_1, \ldots, \alpha_K\}$
    remove the component k_{remove} of the mixture model
    $K_{max} \leftarrow K_{max} - 1$
    $\{\alpha_1, \ldots, \alpha_{K_{max}}\} \leftarrow \left\{ \frac{\widehat{\alpha}_1}{\sum_{k=1}^{k_{max}} \widehat{\alpha}_k}, \ldots, \frac{\widehat{\alpha}_{K_{max}}}{\sum_{k=1}^{k_{max}} \widehat{\alpha}_k} \right\}$
  else
    continue $\leftarrow 0$
  end if
end while
The best solution corresponds to the minimum message length obtained.

5 Data analysis and results

5.1 Evaluating the EM-MML performance on synthetic data

To evaluate the performance of the EM-MML algorithm, we begin by resorting to synthetic data sets: 2-components mixtures of multinomials (10 data sets) and 3-components mixtures (10 data sets) are used. In order to provide useful insights regarding the EM-MML performance, the 10 data sets within each setting exhibit diverse degrees of cluster separation. For this purpose, we measure the separation between all pairs of clusters $(C^k; C^{k'})$ forming in

a partition Π^K according to

$$\text{separation}(\Pi^K) = \frac{2}{k(k-1)} \sum_{k \neq k'} \frac{1}{2} \left[D_{KL}(C^k; C^{k'}) + D_{KL}(C^{k'}; C^k) \right], \quad (10)$$

where

$$D_{KL}(C^k; C^{k'}) = \sum_{l=1}^{L} \sum_{c=1}^{C_l} P(Y_l^k = c) \log \frac{P(Y_l^k = c)}{P(Y_l^{k'} = c)} = \sum_{l=1}^{L} \sum_{c=1}^{C_l} \widehat{\theta}_{klc} \log \frac{\widehat{\theta}_{klc}}{\widehat{\theta}_{k'lc}}$$

is the sum of the Kullback-Leibler divergence between the L multinomial distributions corresponding to components k and k' of the mixture. This symmetric measure of separation ranges from around 0.01 (poor clusters' separation) to 0.17 (good clusters' separation), in the experimental scenarios.

The EM-MML results are compared with those obtained from a standard EM algorithm and well-known information criteria, namely BIC, AIC, CAIC, MAIC, and ICL. The performance of the various criteria and methods is assessed by the rate (over 30 runs) of correct selection of the true number of clusters, and also by the computation time. The results are presented in Figure 2.

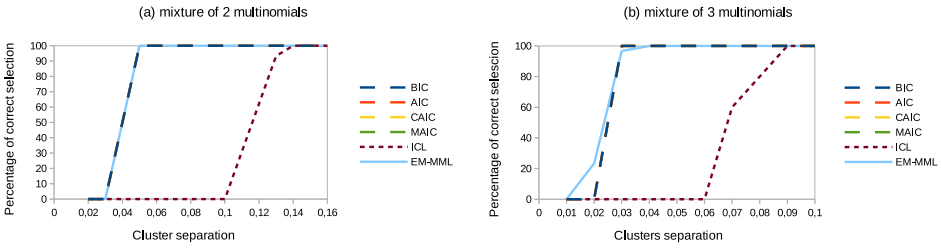

Figure 2: Comparison of several model selection criteria on synthetic data from mixtures of multinomials with $K = 2$ (left) and $K = 3$ (right).

In the two-components data sets with separation lower than 0.04, all the approaches identify only one cluster. ICL is unable to recover the true number of clusters, even with moderate separation (up to around 0.12). The other criteria, and our EM-MML approach, correctly identify two clusters when separation is above 0.04. For three-components mixtures, the ICL criterion is able to identify the correct number of clusters only for cluster separation higher than 0.08. The other criteria (including EM-MML) present similar results identifying three clusters for separations larger than 0.03. Generally one can thus conclude that the EM-MML has a similar performance to BIC, AIC, CAIC, and MAIC in recovering the true number of clusters, while ICL clearly underperforms the other methods for not clearly separated components.

In order to further evaluate the comparative performance of EM-MML, we compare its computation times with that of BIC. Notice that BIC and other information criteria have similar computation times. Based on a 300 runs sample from 2-components mixtures of multinomials, we obtain the following average computation times: 146.84 seconds for EM-MML and 230.67 seconds for BIC. A paired-samples t test yields t = -17.06; df = 299 and p-value < 0.01, allowing to conclude that EM-MML is significantly faster than BIC. For the 3-components mixtures of multinomials, the average of computation times are 194.38 and 239.27 seconds for EM-MML and BIC respectively. Again, the null hypothesis of a paired-samples t test is rejected (t = -6.88; df = 299 and p-value < 0.01), showing that EM-MML is significantly faster than BIC. Overall, the EM-MML method shows good performance when dealing with synthetic data sets: when selecting the true number of clusters it has similar performance to BIC, AIC, CAIC, and MAIC, and outperforms ICL. In terms of computation time, since EM-MML does not require a sequential approach, it becomes clearly faster than the other criteria.

5.2 Experiments on real data

Additional insight into the performance of EM-MML is obtained by applying it to two real data sets from the European Social Survey (ESS). ESS is a biennial survey started in 2002, which measures the attitudes, beliefs, values, and behaviour patterns of European populations. The most recent survey was in 2012 (round 6), covering 30 countries and 243 regions.

For the purpose of our experiment, we aggregate the ESS data by region, taking into account sampling weights (ESS weights) which are meant to provide the sample representativeness. Clustering is performed based on: 1) the *trust in* some institutions (namely, in the country's parliament, legal system, police, politicians, political parties, the European Parliament, the United Nations); 2) the *satisfaction with* life as whole, the economy, the government, and the functioning of the democracy. We recode responses into binary variables: distrust/trust and dissatisfied/satisfied.

The summary of the comparisons of the several model selection criteria on the two ESS data sets is presented in Tables 1 and 2. To measure the relationship between uncovered segments and the clustering base variables, we resort to the Cramer's V association measure - which ranges from 0 (no association) to 1 (perfect association) - and to its sum as a proxy of variables' discriminant capacity.

The number of segments selected by the EM-MML is much lower than for the remaining criteria. This fact avoids estimation problems associated with very small segments and also improves the interpretability of the clustering solution. In addition, the total value of the Cramers V association measure (which evaluates the relationship between clusters and clustering

Table 1: Comparison of Cramer's V association between segmentation base variables "trust in" and the segments obtained with each criterion.

	BIC	AIC	CAIC	MAIC	ICL	EM-MML
Number of segments	11	17	11	18	11	4
Trust in						
country's parliament	.58	.53	.58	.55	.58	.72
the legal system	.55	.54	.55	.57	.55	.67
the police	.52	.51	.52	.50	.52	.57
politicians	.64	.57	.64	.58	.64	.78
political parties	.64	.60	.64	.58	.64	.78
the European Parliament	.46	.45	.46	.47	.46	.52
the United Nations	.47	.43	.47	.48	.47	.52
Sum	3.86	3.63	3.86	3.73	3.86	4.56

Table 2: Comparison of sum of Cramer's V association between segmentation base variables "satisfaction with" and the segments obtained with each criterion.

	BIC	AIC	CAIC	MAIC	ICL	EM-MML
Number of segments	13	18	13	18	13	7
How satisfied with						
life as a whole	.59	.55	.59	.55	.59	.55
present state of economy in country	.58	.55	.58	.55	.58	.62
the national government	.56	.55	.56	.55	.56	.60
the way democracy works in country	.56	.53	.56	.53	.56	.63
Sum	2.29	2.18	2.29	2.18	2.29	2.40

base variables) is higher for EM-MML, indicating variables with higher discriminant capacity in the EM-MML solution. The segments selected by EM-MML criterion are presented in Tables 3 and 4.

Table 3: EM-MML segmentation of the "trust in" data.

Trust in	Segment 1	Segment 2	Segment 3	Segment 4
country's parliament	59.2%	18.2%	14.2 %	35.8%
the legal system	72.0%	23.2%	17.7%	55.4%
the police	81.4%	41.6%	20.8%	72.9%
politicians	45.6%	10.6%	8.7%	20.8%
political parties	44.1%	10.8%	9.9%	19.1%
the European Parliament	41.7%	26.9%	20.2%	28.7%
the United Nations	64.7%	35.0%	24.6%	45.0%

The "trust in institutions" EM-MML segmentation yields four segments. For example, Lisbon is in segment 2, where all the regions have low trust in institutions. The lowest trust values, in this segment, refer to trust in politicians and political parties. Comparing the trust in different institutions, Lisbon citizens have moderate trust in the country police. Dublin,

Table 4: EM-MML segmentation of the "satisfaction with" data.

How satisfied with	Seg. 1	Seg. 2	Seg. 3	Seg. 4	Seg. 5	Seg. 6	Seg. 7
life as a whole	92.5%	93.0%	29.3%	79.7%	71.8%	59.6%	88.0%
present state of economy in country	85.8%	36.4%	3.5%	20.7%	35.5%	15.8%	57.3%
the national government	70.5%	56.7%	8.5%	23.4%	34.5%	21.9%	42.9%
the way democracy works in country	85.8%	81.3%	9.8%	47.1%	46.1%	22.8%	68.2%

another European capital city, is in segment 4, characterized by a higher trust in institutions, specially in the police (72.9%) and the legal system (55.4%). The *separation* of these segments (according to the measure in (10)) is 1.46, indicating that they are well separated.

The seven EM-MML segments for "satisfaction with" are also well separated segments (separation(Π^7) = 1.26). In these segmentation, Lisbon is in segment 6 and Dublin in segment 4. In both capitals, people only feel satisfied with life as a whole, but in Dublin they feel generally more satisfied than in Lisbon.

6 Discussion and perspectives

In this work, a model selection criterion and method for finite mixture models of categorical observations was proposed. The new algorithm simultaneously performs model estimation and selects the number of components/clusters. When compared to information criteria, which are commonly associated with the use of the EM algorithm, the EM-MML method exhibits several advantages: 1) it easily recovers the true number of clusters in synthetic data sets with various degrees of separation; 2) its computations times are significantly lower than those required by the sequential use of EM in standard approaches (such as BIC); 3) when applied to real data sets it produces more parsimonious solutions, thus easier to interpret. An additional advantage of the proposed approach that stems from obtaining more parsimonious solutions is that such solutions have a higher number of observations per cluster, thus helping to overcome eventual estimation problems.

Since the performance of the EM-MML is encouraging for selecting the number of clusters, and the same criterion was already used for feature selection [8], future developments will include the integration of both model and feature selection.

References

[1] Chen T, Zhang NL, Liu T, Poon KM, Wang Y. Model-based multidimensional clustering of categorical data. Artificial Intelligence. 2012; 176:22462269.

[2] Xiong T, Wang S, Mayers A, Monga E. Dhcc: Divisive hierarchical clustering of categorical data. Data Mining and Knowledge. 2012;24:103–35.

[3] Giordana M, Diana G. A clustering method for categorical ordinal data. Communications in Statistics-Theory and Methods. 2011;40:13151334.

[4] Chen K, Liu L. Best k: Critical clustering structures in categorical datasets. Knowledge and Information Systems. 2009;20:1–33.

[5] Desai A, Singh H, Pudi V. Disc: Data-intensive similarity measure for categorical data. Advances in Knowledge Discovery and Data Mining, Lecture Notes in Computer Science. 2011;6635:469–481.

[6] Constantinopoulos C, Titsias MK, Lika A. Bayesian feature and model selection for gaussian mixture models. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2006;28:1013–1018.

[7] Celeux G, Martin O, Lavergne C. Mixture of linear mixed models for clustering gene expression profiles from repeated microarray experiments. Statistical Modelling. 2005;5:243–267.

[8] Law M, Figueiredo M, Jain A. Simultaneous feature selection and clustering using mixture models. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2004;26:1154–1166.

[9] Bouguila N. Clustering of count data using generalized dirichlet multinomial distributions. IEEE Transactions on Knowledge and Data Engineering. 2008;20:462–474.

[10] Li M, Zhang L. Multinomial mixture model with feature selection for text clustering. Knowledge-Based Systems. 2008;21:704–708.

[11] Deya T, Lim CY. Comparisons of computational methods for clustered binary data. Journal of Statistical Computation and Simulation. 2013; 83:20302046.

[12] Schwarz G. Estimating the dimension of a model. The Annals of Statistics. 1978;6:461–464.

[13] Akaike H. Maximum likelihood identification of gaussian autorregressive moving average models. Biometrika. 1973;60:255–265.

[14] Dempster A, Laird N, Rubin D. Maximum likelihood estimation from incomplete data via the EM algorithm. Journal of Royal Statistical Society. 1997;39:1–38; series B.

[15] Figueiredo MAT, Jain AK. Unsupervised learning of finite mixture models. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2002;24:381–396.

[16] Silvestre C, Cardoso M, Figueiredo M. Clustering with finite mixture models and categorical variables. In: International Conference on Computational Statistics - COMPSTAT2008; 2008.

[17] Everitt BS, Hand DJ. Finite mixture distributions. New York: Chapman and Hall; 1981.

[18] Titterington DM, Smith AFM, Makov UE. The statistical analysis of finite mixture models. New York: Wiley; 1985.

[19] McLachlan GJ, Peel D. Finite mixture models. New York: Wiley; 2000.

[20] Titterington DM, Smith AFM, Makov U. Statistical analysis of finite mixture distributions. Jon Wiley and Sons; 1985.

[21] Melnykov V, Maitra R. Finite mixture models and model-based clustering. Statistics Surveys. 2010;4:80116.

[22] Lindsay BG. Mixture models: Theory, geometry, and applications. Hayward: Institute of Mathematical Statistics; 1995.

[23] Wedel M, Kamakura WA. Market segmentation: Conceptual and methodological foundations (international series in quantitative marketing). 2nd ed. Massachusetts: Kluwer Academic Publishers; 2000.

[24] Newcomb S. A generalized theory of the combination of observations so as to obtain the best result. American Journal of Mathematics. 1886; 8:343–366.

[25] Pearson K. Contributions to the mathematical theory of evolution. Philosophical Transactions of the Royal Society. 1894;185:71–110.

[26] Macdonald P, Pitcher T. Age-groups from size-frequency data: a versatile and efficient method of analysing distribution mixtures. Journal of the Fisheries Research Board of Canada. 1979;36:987–1001.

[27] Do K, McLachlan GJ. Estimatiom of mixing proportions: A case study. Applied Statistics. 1984;33:134–140.

[28] Al-Hussaini EK, Abdel-Hamid AH. Accelerated life tests under finite mixture models. Journal of Statistical Computation and Simulation. 2006;76:673–690.

[29] Maitra R, Melnykov V. Finite mixture models and model-based clustering. Statistical Surveys. 2010;4:80–116.

[30] Elmore RT, Wang S. Identifiability and estimation in finite mixture models with multinomial components. Technical Report 03-04, Department of Statistics, Pennsylvania State University. 2003;.

[31] Portela J. Clustering discrete data through the multinomial mixture model. Communications in Statistics-Theory and Methods. 2008; 37:3250–3263.

[32] Vermunt JK, Magidson J. Latent class cluster analysis. Applied latent class analysis. 2002;11:89–106.

[33] McLachlan G, Thriyambakam K. The EM algorithm and extensions. 2nd ed. New Jersey: John Wiley and Sons; 2008.

[34] Gupta M, Chen Y. Theory and use of the EM algorithm. Now Publishers; 2011.

[35] Celeux G, Martin-Magniette ML, Maugis-Rabusseau C, Raftery AE. Comparing model selection and regularization approaches to variable selection in model-based clustering. Journal de la Société Française de Statistique. 2014;155:57–71.

[36] Fraleya C, Raftery AE. Model-based clustering, discriminant analysis, and density estimation. Journal of the American Statistical Association. 2002;97:611–631.

[37] Bozdogan H. Model selection and akaike's information criterion (AIC): The general theory and its analytical extensions. Psycometrika. 1987; 52:345–370.

[38] Bozdogan H. Mixture-model cluster analysis using model selection criteria and a new informational measure of complexity. In: Proceedings of the First US/Japan Conference on the Frontiers of Statistical Modeling: An Approach; 1994. p. 69–113.

[39] Biernacki C, Celeux G, Govaert G. Assessing a mixture model for clustering with the integrated completed likelihood. IEEE Transactions on Pattern analysis and Machine Intelligence. 2000;22:719–725.

[40] Rissanen J. Complexity in statistical inquiry. World Scientific; 1989.

[41] Wallace C, Boulton D. An information measure for classification. The Computer Journal. 1968;11:195–209.

[42] Cover T, Thomas J. Entropy, relative entropy and mutual information. Elements of Information Theory; 1991. Chapter 2.

[43] Baxter RA, Olivier JJ. Finding overlapping components with mml. Statistics and Computing. 2000;10:5–16.